

HPC Systems Overview: Blue Gene /Q Systems, Storage, Software, and Other Notes

William Scullin
Operations Group, LCF



Building Blocks of the Universe

The Periodic Table of Elementary Particles and Forces

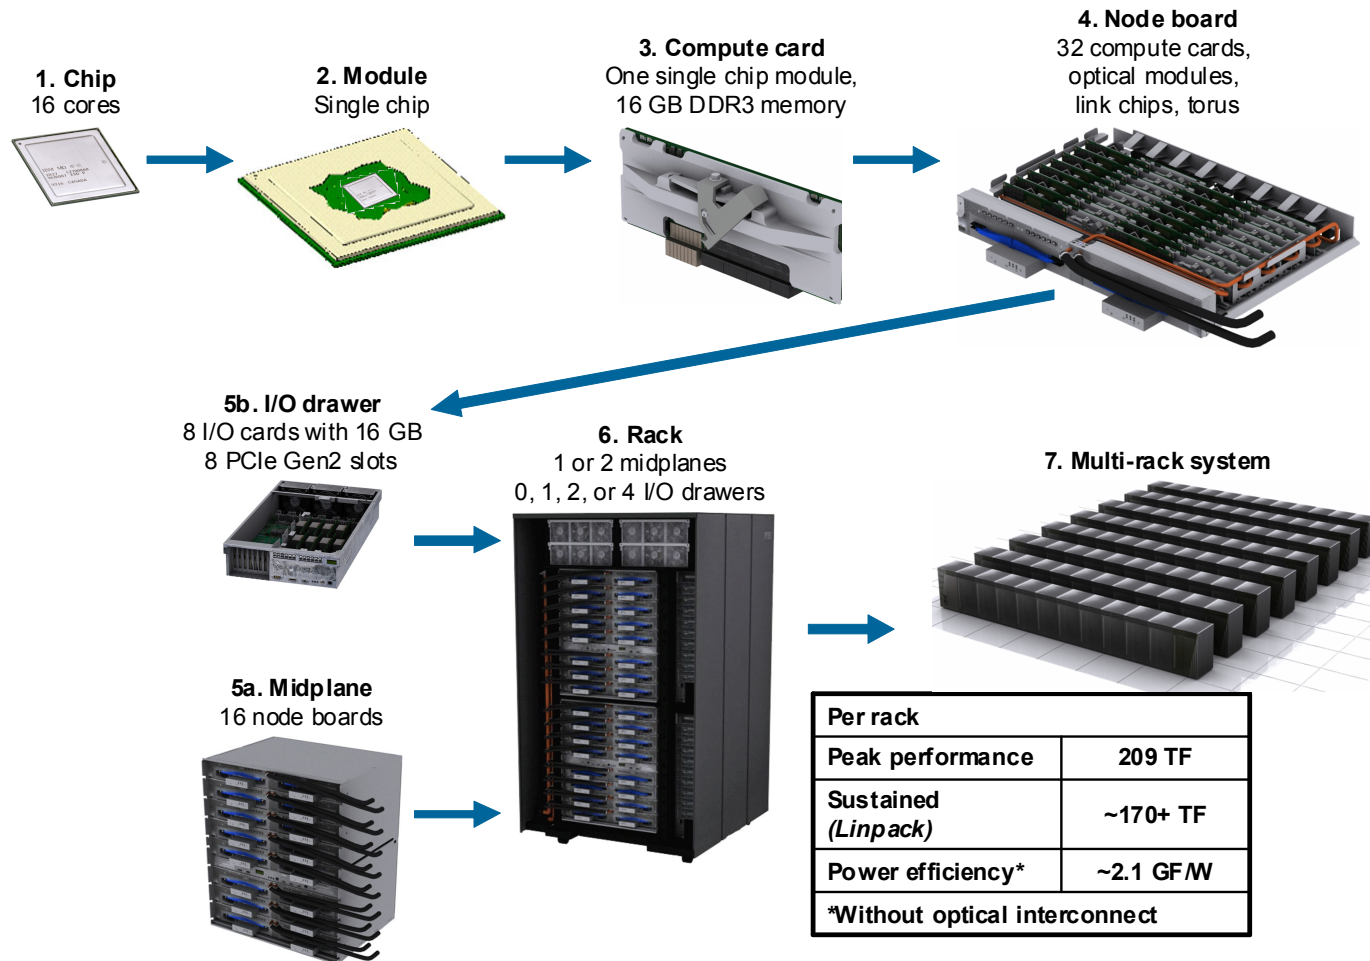
Three Generations of Matter (Fermions)

	I	II	III	
mass	2.4 MeV	1.27 GeV	171.2 GeV	0
charge	$\frac{2}{3}$	$\frac{2}{3}$	$\frac{2}{3}$	0
spin	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1
name	u up	c charm	t top (truth)	γ photon (electromagnetic)
	4.8 MeV	1.04 MeV	4.2 GeV	0
	$-\frac{1}{3}$	$-\frac{1}{3}$	$-\frac{1}{3}$	0
	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1
Quarks	d down	s strange	b bottom (beauty)	g gluon (strong force)
	<2.2 eV	<0.17 MeV	15.5 MeV	91.2 GeV
	0	0	0	0
	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1
	ν_e electron neutrino	ν_μ muon neutrino	ν_τ tau neutrino	Z weak force
	0.511 MeV	105.7 MeV	1.777 GeV	80.4 GeV
	-1	-1	-1	±1
	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1
Leptons	e electron	μ muon	τ tau	W weak force
				125-186 GeV
				±1
				0
				H higgs boson

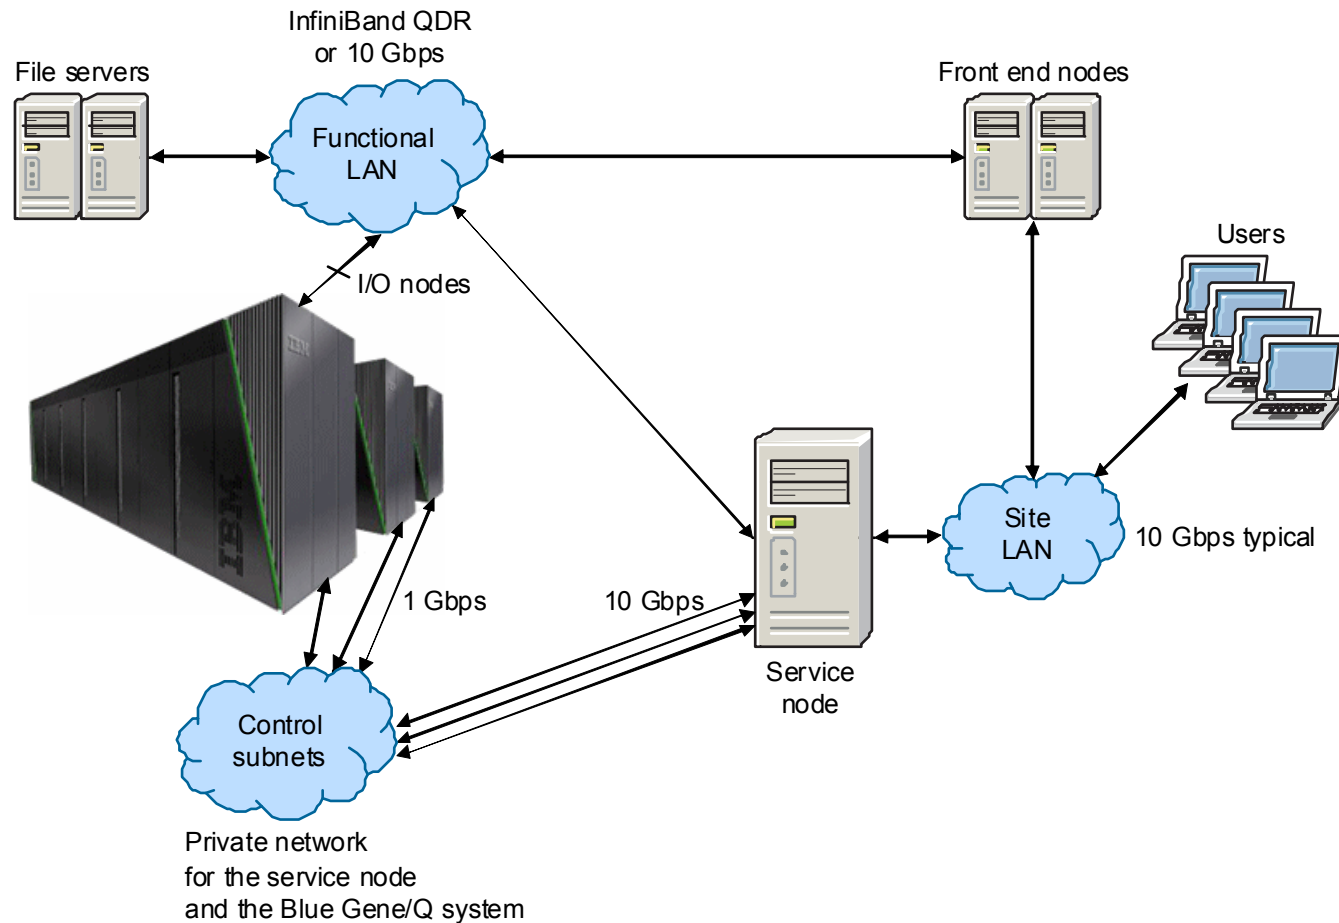
Bosons (Forces)



Anatomy of a Blue Gene /Q (Not Ours)



No Blue Gene /Q is an island



System Details

Blue Gene /Q System	Vesta	Mira	Cetus
Function	Test & Development	INCITE Production	INCITE Test & Development
Login address	vesta.alcf.anl.gov	mira.alcf.anl.gov	cetus.alcf.anl.gov
Login OS	RHEL 6.3	RHEL 6.3	RHEL 6.3
Login CPUs	1 Power 740: - Power 7 - 48 cores @ 3.7 GHz	4 Power 740(s): - Power 7 - 48 cores @ 3.7 GHz	1 Power 740: - Power 7 - 48 cores @ 3.7 GHz
Login memory	64GB	64GB	64GB
BG /Q CPU (16 core)	1.6 GHz PowerPC A2	1.6 GHz PowerPC A2	1.6 GHz PowerPC A2
# Nodes / # Cores	2048 / 32,768	49,152 / 786,432	1024 / 16,384
BG /Q Memory	32TB (16GB per node)	640TB (16GB per node)	16TB (2GB per node)
# I/O nodes (Ratio)	64 @ IB (32:1)	384 @ IB (128:1)	8 @ IB (128:1)
Blue Gene /Q Driver	V1R2M0 + efixes to 30	V1R2M0 + efixes to 27	V1R2M0 + efixes to 27

- Login to ALCF Resources is via ssh with cryptocard authentication
- Round robin DNS will place you on an available login node named:
 <machine>lac{number}.alcf.anl.gov
- Logins are for compilation and job submission only -
- You may use parallel make, but be gentle



System Details

Analytics System	Tukey
Function	INCITE Production
Login address	tukey.alcf.anl.gov
Node OS	RHEL 6.3
Node CPUs	Two 2GHz AMD Opteron 6128 CPUs per node (8 cores per CPU, 16 cores total)
Node memory	64GB
GPUs	Two NVIDIA Tesla M2070 GPUs per node
GPU memory	6GB
Nodes	100
Interconnect	QDR Infiniband interconnect
Local scratch	450 GB

- Login to ALCF Resources is via ssh with cryptocard authentication
- Round robin DNS will place you on an available login node named:
login{number}.tukey
- Logins are for compilation and job submission only
- Four tukey nodes are currently in use for GPFS token management



Driver and OS status

- ▶ We're currently running Red Hat Enterprise Linux 6.3
 - this covers logins, the control system and infrastructure, and IO nodes
 - We are continuously patching where possible
 - holding back on anything that strongly impacts the toolchain, control system, or availability
 - 99% of packages are stock RHEL 6.3 ppc64 packages
 - everything is 64-bit
 - we will install 32-bit packages and applications if absolutely necessary
 - so far there haven't been any cross-compiling surprises
- ▶ Current driver is the V1R2M0 driver
 - Vesta (testing and development) receives all efixes first with about two weeks between an efix being installed on Vesta and installation in the production environment
 - Vesta currently is on efix 30
 - Cetus and Mira are not yet - efix 30 has problems
 - We're expecting updated sources... soon.



- ▶ Be a good citizen!
 - Don't run parallel make for more than a fourth of the cores (16) than a login has
 - Clean up your files when done
 - We purge /tmp, but it's better if you clean up after compiles
 - We will terminate processes impacting stability and other's use of the logins
 - We will contact you if we do
 - We can usually accommodate your workflows – just ask
- ▶ Cetus and Mira
 - All important filesystems are synced and shared
 - Cetus is for testing and development jobs - not production
 - You cannot submit Cetus (prod) jobs from Mira (prod-devel) and vice-versa
- ▶ Vesta file systems are not shared on Cetus and Mira for security reasons
- ▶ If you lack a .soft file file at login we can create one or basic content should look like:

@default
+mpiwrapper-xl

Filesystems and Layout

- /soft
 - we'll get back to this
 - mounted everywhere
 - a GPFS bind mount
- /projects
 - mounted everywhere, not backed up
 - check quota with /usr/local/bin/myprojectquotas
 - GPFS
 - also referred as
 - /gpfs/mira-fs0/projects on Mira, Cetus, and Tukey
 - /gpfs/vesta-fs0/projects on Vesta
- /home
 - mounted everywhere, backed up on Mira, Cetus, and Tukey only
 - check quota with /usr/local/bin/myquota
 - GPFS
 - also referred to as
 - /gpfs/mira-home on Mira, Cetus, and Tukey
 - /gpfs/vesta-home on Vesta



/soft layout and finding things



- ▶ We're trying to reorganize /soft to make things easier to follow
 - Arrangements will be by function, ie: compilers in /soft/compilers, performance tools in /soft/perftools, softenv and modules in /soft/environment.
- ▶ softenv keys should be authoritative
- ▶ front end software (editors, X, games) is installed in RHEL's default locations



Cetus



Vesta



Mira



Tukey (or part of it)



A Random Mantis Shrimp



Cobalt

- Very similar to /P with caveats
 - modes are different
 - `-c{1,2,4,8,16,32,64}` sets ranks per core
 - `-n` gives you nodes
 - `--proccount` gives you total processes
 - custom IO kernels are not yet supported, but coming soon
 - the `cqsub` and `cqstat` commands are going the way of the dodo
 - use `qstat` and `qsub`
- Script mode is the same
 - block starts off booted
 - `cobalt-subrun` does not work
 - see wiki for details (use `runjob` in scripts and consult the man page - only in scripts)

Block naming and Cobalt

- Block names follow logical names, not hardware names
- Why?
 - 32 character limit on block names
 - Mira allows many more degrees of freedom in block configurations
 - Allows us to state which midplanes, and by extension which hardware is in use in a given location when the hardware locations make little sense
 - Makes sub-block setup easier
- One rack has the topology 4x4x4x8x2
- One midplane is 4x4x4x4x2



Decoding Block Names

- ▶ LOC-CCCC-XXXX-[T]-[PPPP]-SIZE
- ▶ LOC = location identifier, like ANL, CHR, VES, CET, MIR, EAS can be up to 7 characters.
- ▶ CCCC = The bottom right front corner as described as a set of 5-dimensional coordinates ABCDE. This corresponds to the node location of the node of rank 0 in a ABCDE-type mapping scheme (node 0).
- ▶ XXXX = The top left rear corner of the block in each dimension (node n-1).
- ▶ T = an optional identifier indicating which dimensions are Mesh and which are torus. This is a bitmasked value (0 = torus, 1 = mesh). No value implies the maximum number of torus dimensions for that block
- ▶ PPPP = indicator of passthrough extents in each dimension. This will have a value of 0, 1, or 2.
- ▶ SIZE = The overall size in nodes of the block. This should correspond to the product of the extents.



Block name example

- ▶ A sample logical address could be: MIR-04C00-48FF2-7-2048
- ▶ Think LOC-CCCCC-XXXXX-[T]-[PPPP]-SIZE
- ▶ This corresponds to:
 - one midplane in the A dimension, first midplane in A
 - one midplane in the B dimension, starting at the second midplane (row 1, to be exact)
 - one midplane in the C dimension,
 - Four midplanes in the D dimension,
 - A,B,C dimensions are mesh, D is a torus
 - 2048 Nodes
- ▶ Old style it might be: MIR-R14-R15-2048



Rules of the Road

- ▶ Like /P, on /Q blocks must have a sufficient ratio of IO nodes to Compute Nodes
- ▶ Unlike /P, /Q IO blocks are persistent and may be shared
 - writing to /tmp and not removing files is anti-social
 - we do not offer a way to retrieve files written to /tmp at this time
 - crashing an IO block crashes compute blocks reliant on that IO block
- ▶ All GPFS filesystems available on the login nodes are available on the IO nodes for access from the computes
- ▶ You are always the only thing on the compute node's compute cores and memory

Resource Isolation (Mira and Cetus)

- ▶ Resource sets where you are the only one on our resources.
 - ≥ 512 (midplane or above): The IONs, computes and blocks are all yours within that block)
- ▶ You are shared when:
 - ≥ 256 : you share IONs
 - > 128 : you share the block with other users (i.e. 2 64s may have different users). May see traffic on interconnect from other users (I think, I figure this has to be the case for traffic to the ION through J06 and J11)
- ▶ You are always the only thing on the compute node's compute cores and memory

Resource Isolation (Vesta)

- Vesta is configured for users doing OS development
- Resource sets where you are the only one on our resources.
 - ≥ 128 (4 node boards): The IONs, computes and blocks are all yours within that block)
- You are shared when:
 - > 128 : you share IONs
 - > 32 : you share the block with other users
- You are always the only thing on the compute node's compute cores and memory
- Crashing an IO block takes out jobs on any attached compute blocks

- Every user must have at least one Project they are assigned to
- Projects are then given allocations
 - Allocations have an amount, start, and end date and are tracked separately; Charges will cross allocations automatically. The allocation with the earliest end date will be charged first, until it runs out, then the next, and so on
- Charges are based on the partition size, NOT the # of nodes or cores used!
- Reservations are charged for the full time they are active
- Will be managed with clusterbank
 - Use the 'cbank' command (see 'cbank --help')
- Examples:
 - # list all charges against a particular project
 - `cbank -l charge -p <projectname>`
 - # list all active allocations for a particular project
 - `cbank -l allocation -p <projectname>`



When things go wrong... running

- Cobalt jobs, by default, produce three files prefixed with either the job number or the name specified with qsub's **-O** option:
 - **\$PREFIX.output** with output to standard out from you application
 - **\$PREFIX.error** with output from the control system, scheduler, and your application directed to standard error
 - **\$PREFIX.cobaltlog** with a record of the environment and the submission command
 - Generated at submit time
 - Very useful for the support team in debugging
- Only cobaltlog is generated at submit time, the others at runtime
- At boot, the .error file will have a non-zero size for non-script jobs
 - Most of the messages are related to booting, it's a decent way to follow startup progress
 - Script jobs leave the handling of standard error up to you, though most do use the .error file
 - We'll walk through a normal boot in a moment...
- If you think there is an issue, it's best to save all three files
- If you have a script job, save the script in an unedited state

When things go wrong... running

- You'll see RAS events appear in your .error file it's not always the sign of trouble
- Exit codes are more meaningful
 - 0 is normal
 - 15 is usually hitting walltime and dying with SIGTERM
 - 9 is usually hitting walltime and requiring SIGKILL, but may be a system issue
 - anything else - unless running Nek - is worth reporting
- RAS stands for Reliability, Availability, and Serviceability
- Few are a sign of a serious issue, most are system noise
 - Messages have a severity associated with them
 - INFO
 - WARN
 - ERROR
 - FATAL
 - Only FATAL RAS events will lead to the termination of your application
 - Still worth watching as they may be the sign of an application issue





Remember:
The VelociRAStor says **ONLY FATAL RAS EVENTS ARE FATAL**

Thanks for listening!
Any questions?

